

FR3D: Three-dimensional flow reconstruction and force estimation for unsteady flows around extruded bluff bodies via conformal mapping aided convolutional autoencoders

Ali Girayhan Özbay*, Sylvain Laizet

Department of Aeronautics, Imperial College London, Exhibition Road, London, SW7 2AZ, United Kingdom

ARTICLE INFO

MSC:
68T07
76F65

Keywords:
Flow reconstruction
Neural networks

ABSTRACT

In many practical fluid dynamics experiments, measuring variables such as velocity and pressure is possible only at a limited number of sensor locations, for a few two-dimensional planes, or for a small 3D domain in the flow. However, knowledge of the full fields is necessary to understand the dynamics of many flows. Deep learning reconstruction of full flow fields from sparse measurements has recently garnered significant research interest, as a way of overcoming this limitation. This task is referred to as the flow reconstruction (FR) task. In the present study, we propose a convolutional autoencoder based neural network model, dubbed FR3D, which enables FR to be carried out for three-dimensional flows around extruded 3D objects with different cross-sections. An innovative mapping approach, whereby multiple fluid domains are mapped to an annulus, enables FR3D to generalize its performance to objects not encountered during training. We conclusively demonstrate this generalization capability using a dataset composed of 80 training and 20 testing geometries, all randomly generated. We show that the FR3D model reconstructs pressure and velocity components with a few percentage points of error. Additionally, using these predictions, we accurately estimate the Q-criterion fields as well lift and drag forces on the geometries.

1. Introduction and related work

Flow reconstruction (FR) involves the prediction of dense fields such as velocity based on sparse measurements. Since typical experiments in fluids involve only point measurements of the flow via simple and inexpensive methods such as pitot tubes, FR techniques can provide researchers additional insight into flows when more advanced techniques such as particle image velocimetry (PIV) are not available.

Various statistical tools have been applied to FR such as linear stochastic estimation (LSE) (Bonnet et al., 1994), gappy proper orthogonal decomposition (gappy POD) (Willcox, 2006), extended proper orthogonal decomposition (EPOD) (Borée, 2003), and sparse representation (Callaham et al., 2019). Though these techniques are time-tested and have been applied in practical experiments, for instance to estimate and control the flow in a backward-facing step case via LSE (Taylor and Glauser, 2004), their linear nature limit their capability to deal with complex flows.

Neural networks (NNs), owing to their universal approximation capabilities (Funahashi, 1989), are capable of learning arbitrary non-linear and high-dimensional relationships in datasets. This capability

makes them very attractive for FR tasks. As a result, the recent explosion of interest in NNs – enabled by substantial increases in computing power, theoretical advances, and the availability of open-source deep learning software – has coincided with a shift towards NN-based FR, and substantial strides were made with the application of NNs to the field. Notably, Erichson et al. (2020) produced a seminal study exploring the usage of neural networks to reconstruct flows past cylinders. A number of works followed Erichson et al. a selection of which are presented: Fukami et al. (2020) demonstrated that NN-based methods can outperform linear FR methods for the reconstruction of flows past cylinders and flapped airfoils, and also coupled NN-based FR with Voronoi tessellations to achieve flexibility in terms of the sensor setup (Fukami et al., 2021). Sun and Wang (2020) investigated the application of physics-informed Bayesian NNs in FR, demonstrating high robustness to noise when reconstructing flows in simulated vascular structures. Dubois et al. (2022) extended FR to variational autoencoder architectures. Kumar et al. (2021) developed a recurrent NN architecture to carry out FR with extremely sparse sensor setups.

* Corresponding author.

E-mail address: aligirayhan.ozbay14@imperial.ac.uk (A.G. Özbay).

Xu et al. (2023) considered the usage of physics-informed loss functions to train FR models using gappy data. He et al. (2022) explored the usage of graph attention NNs to reconstruct flows past cylinders. Carter et al. (2021) investigated the usage of an NN-based FR model on experimental data of flow past an airfoil. Applications outside typical wind tunnel-like scenarios, such as in porous media (Akeweje et al., 2023) or the reconstruction of atmospheric flows based on satellite imagery (Schweri et al., 2021), have also been developed.

However, the applicability of even these recent approaches to practical scenarios is limited. The first obstacle is tackling multi-geometry FR. Even recently published works typically investigate FR for a single geometry only (often a 2D circular cylinder), and as a result the models used in such works must be re-trained for every case investigated; a model trained on e.g. a circular cylinder will not work well for a square cylinder. This necessitates laborious data collection and a computationally expensive training process. To overcome this limitation, a growing body of works have investigated 2D multi-geometry FR using techniques such as graph convolutional neural networks (Chen et al., 2021; Liu et al., 2022; Duthé et al., 2023) and conformal mappings (Özbay and Laizet, 2022a).

The second challenge pertains to the reconstruction of three-dimensional flows, regarding which relatively few works exist compared to the reconstruction of two-dimensional flows. Three dimensional flows exhibit substantially more complicated dynamics than two-dimensional flows (Verma et al., 2018), and require much greater computing power to process. Despite this, a growing body of works is tackling the challenge of 3D FR. Particularly, reconstruction of 3D flows past cylinders (Laima et al., 2023) and of flows concerning domains without embedded objects such as channel flows (Güemes et al., 2021; Guastoni et al., 2021) have been investigated recently. A number of studies have also investigated the reconstruction of 2D slices of flows past square cylinders (Nakamura and Fukagata, 2022; Matsuo et al., 2021) and vice versa (Pérez et al., 2020) (i.e. reconstruction of 3D fields from 2D slices).

In this work, we introduce a method enabling the reconstruction of unsteady three-dimensional flows around objects with arbitrarily shaped cross-sections. This is achieved via an autoencoder-based convolutional neural network architecture which incorporates conformal mappings to achieve geometry invariance (Özbay and Laizet, 2022a,b). In our previous study (Özbay and Laizet, 2022a,b), we have shown that it is possible to reconstruct dense contemporaneous or future vorticity fields of two-dimensional flows past various objects from current sparse sensor measurements, even for objects not encountered during training. To achieve optimal performance in these tasks, Schwarz–Christoffel mappings were used for choosing the sampling points of the dense fields. The results showed that the mapping aided approach provides a substantial boost in accuracy for all model and sensor setup configurations, enabling percentage errors under 3%, 10% and 30% for reconstructions of pressure, velocity and vorticity fields, respectively.

For the present study, we reconstruct unsteady three-dimensional flows around bluff bodies with periodic spanwise boundary conditions at $Re = 500$ (based on the freestream velocity and the characteristic length of the bluff body). Additionally, we use these reconstructions to accurately predict the aerodynamic forces experienced by the investigated geometries as well as the Q-criterion (Jeong and Hussain, 1995), which defines vortices as areas where the vorticity magnitude is greater than the magnitude of the rate of strain.

The paper is organized as follows: first, in Section 2, we detail the procedure used to generate our dataset and the experiments to be carried out using the dataset. Next, in Section 3, we expound upon the FR3D model architecture and its training procedure. Subsequently, in Section 4, we display the performance of the FR3D model. Finally, in Section 5, we summarize the results, and identify avenues for further research in 3D flow reconstruction.

2. Data and experimental setup

2.1. Geometries, meshing and flow simulations

Our dataset consists of 100 geometries $G_i, i \in [0, 99]$, randomly generated using a method based on Bezier curves by Viquerat and Hachem (2020). Each geometry uses 4 control points for the curves, chosen randomly in a square domain with characteristic length L_m , which enables the generation of convex as well as concave shapes. Fig. 1 showcases the diversity of the geometries created in this manner, including airfoil-like cross-sections, objects with concavities and objects with sharp corners.

Each geometry G_i was placed in a $40L_m/7 \times 40L_m/7$ square domain, extruded for $20L_m/7$ in the spanwise direction. The domains were meshed using an automated procedure with c. 30,000 hexahedral and triangular prism elements each, with wake refinement applied. The flows were computed at a Reynolds number equal to $Re = u_\infty L_m/\nu = 500$ using the PyFR solver (Witherden et al., 2014) (u_∞ is the free stream velocity and ν is the viscosity of the flow). This solver is a flux reconstruction (Huynh, 2007) based advection–diffusion equation solver using the artificial compressibility approach to solve the incompressible Navier–Stokes equations. It was chosen for its Python interface and GPU acceleration capabilities. 800 snapshots between $\tau^* = 5.71$ and $\tau^* = 17.14$ (when the flow is fully established) were recorded per geometry, where $\tau^* = u_\infty \tau/L_m$ is the time normalized by the large eddy turnover time L_m/u_∞ .

The snapshots collected were split into training and validation datasets, with all snapshots belonging to 20 randomly chosen geometries constituting the validation set, and the rest serving as the training set. This setup ensures that our model must have reasonable generalization performance to perform well on the validation set, as learning the reduced-order dynamics of specific flows (as in e.g. the Dynamic Mode Decomposition Schmid, 2010) is not sufficient to reconstruct flows past unseen geometries.

2.2. Flow validation

The solver settings used to generate the flow dataset were validated using the canonical case of the flow past a cylinder with diameter D in a spanwise periodic domain. Two quantities were compared against reference data by Mittal et al. (1997) at $Re = 300$, both averaged in the spanwise direction:

- $\langle u \rangle - u_\infty$: Time averaged u -velocity deficit
- $\langle u'u' \rangle / u_\infty^2$: Streamwise-component of the Reynolds stress tensor

Our results were obtained using an automatically generated mesh, created with the same procedure as the one used to generate the meshes for the random geometries. The solver settings were kept identical to the settings used to obtain flow solutions for the random geometries, save for adjusting u_∞ to obtain the correct Reynolds number. The solution took approximately 1 h to complete on a single Nvidia A100 GPU. The two quantities are plotted at several downstream locations in Fig. 2.

The plots show good agreement between the reference values and our flow solution. u -velocity deficit profiles, including the peak deficit, are replicated with minimal errors at all three downstream locations. Streamwise Reynolds stress tensor results also show good agreement throughout most of the domain – important features such as the dual peaks of the profile are closely followed – although the dip between the two peaks is underestimated.

As a further soundness check, we also compared the time averaged lift coefficient (C_L) and drag coefficient (C_D) values obtained from our simulation with published data. We recorded a C_D value of 1.46, which is in line with the range of values between 1.22 and 1.50 in previous literature as compiled by Giannenas and Laizet (2021), and a




Fig. 1. Surface meshes of a selection of the randomly generated geometries used in the study.

C_L of 2.8×10^{-5} – i.e. practically 0 for a single precision calculation – as expected. Thus, as our meshing and solver settings produce results that are largely in line with previous literature for this validation case, we can be confident that our dataset consists of physically correct snapshots.

2.3. Postprocessing

As the final step in our data generation process, postprocessing was applied to the snapshots to obtain sets of inputs and outputs used to train the neural network models. Our postprocessing extends the methodology which was used in our previous works based on conformal mappings to incorporate geometry invariance in neural network based FR methods (Özbay and Laizet, 2022a,b).

In the aforementioned method, the fluid domain F_i around each geometry G_i is treated as a doubly connected 2D region, and a mapping f between an annulus and F_i is computed, as visualized in Fig. 3. Subsequently, a grid equispaced in the radial and angular dimensions is generated in the annular domain, and the mapping is used to calculate the coordinates of the gridpoints in F_i , which constitute the sampling points of the ground truth dense fields. As a result, each slice of the grid along the radial direction is guaranteed to start on the surface of the geometry G_i and end on the outer boundary of the fluid domain.

This work extends the method by incorporating the extrusion of the 2D cross section in the third dimension. In essence, this allows the dense flow fields around each geometry to be represented in cylindrical coordinates (r, θ, z) . We choose a resolution of 64 grid points per dimension; hence, the dense fields are sampled on a $64 \times 64 \times 64$ grid.

We consider two experiments with different types of sensor inputs. In the first experiment, the sensor inputs consist of pressure and velocity probes. The former consist of 50 probe locations equispaced in the angular direction along the inner ring in the annular domain at 10 stations equispaced in the spanwise direction, for a total of 500 pressure sensors. The latter, also consisting of 500 sensors, are arranged in a $10 \times 10 \times 5$ grid, spanning a $L_m \times 5L_m/3 \times 20L_m/7$ box $L_m/7$ downstream of the trailing edge of each object. Using this setup, dubbed the “sparse” setup, we consider the reconstruction of the four primary variables in the Navier–Stokes equations — the pressure p and the three velocity components u, v and w .

The second experiment retains the previous pressure sensor setup but uses plane measurements of velocity fields, called the “plane” setup. Two perpendicular planes with 32×32 velocity sensors each are considered; an xy -plane and an xz -plane, both downstream of the randomly generated objects. This sensor setup has seen use in experiments (Chandramouli et al., 2019), and hence is of interest within the context of certain setups involving PIV, an optical technique used

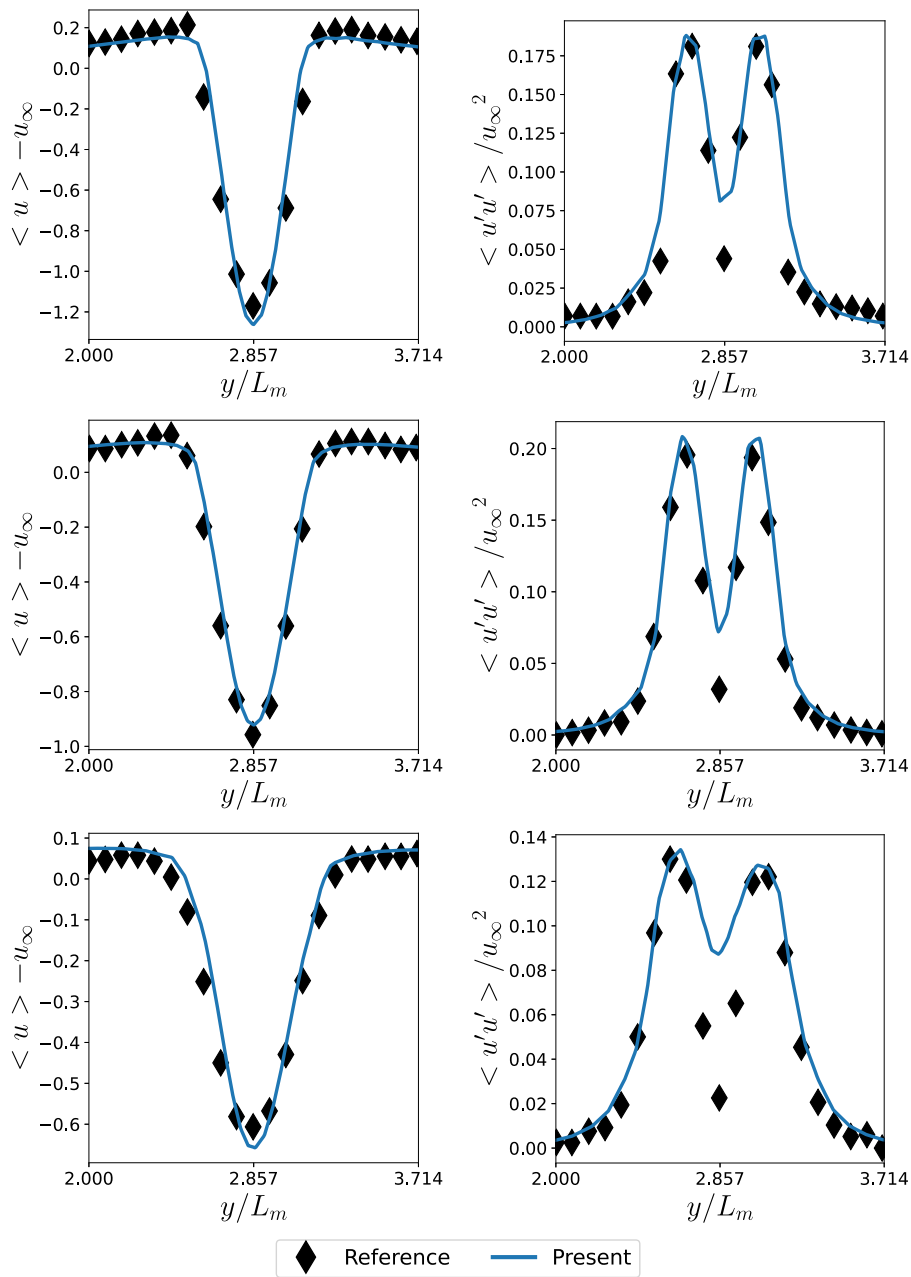


Fig. 2. Comparison of results by Mittal et al. (1997) with our results for $\langle u \rangle - u_\infty$ (left) and $\langle u'u' \rangle / u_\infty^2$ (right) for the $Re = 300$ cylinder, 1.2 D (top), 1.5 D (middle) and 2.0 D (bottom) downstream the center of the cylinder.




Fig. 3. A random geometry (left) and its corresponding preimage (right). Blue and green contours depict the norm and argument in the annular domain, respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

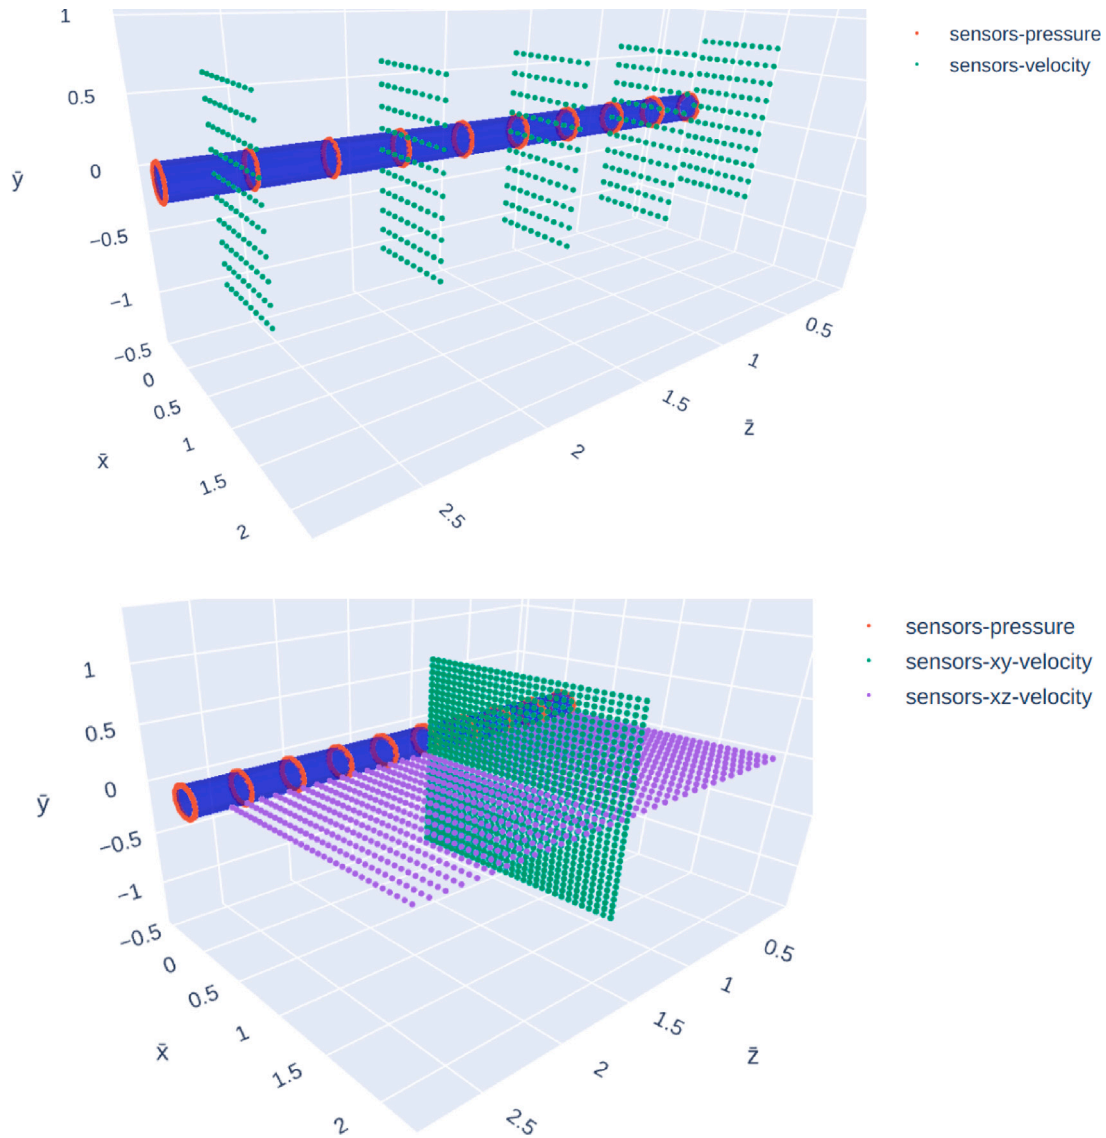


Fig. 4. 3D representation of the sparse (top) and plane (bottom) sensor setups on a circular cylinder. $\bar{x}, \bar{y}, \bar{z}$ refer to the x, y, z coordinates normalized by L_m .

in experiments capable of measuring the velocity field for turbulent fluid flows, and for which the pressure fields must be inferred from the velocity measurements (De Kat and Van Oudheusden, 2012). The two sensor setups are visualized in Fig. 4.

We would like to point out that our flow reconstruction framework has been designed to be trained with three-dimensional high-fidelity computational data. Experimental three-dimensional data (such as 3D PIV) could also potentially be used in the future for the training of our framework.

3. Model architecture and training

To carry out 3D flow reconstruction, we use a convolutional autoencoder-based architecture, dubbed FR3D. It consists of three parts: an encoder, a decoder and a latent embedder. The encoder

$$\mathcal{E} : \mathbb{R}^I \rightarrow \mathbb{R}^E$$

compresses its input \mathbf{x} into a latent space embedding \mathbf{l} , where $I \ll E$. Its architecture is a typical convolutional encoder architecture; each block consists of an initial convolution to double the number of channels, followed by sub-blocks applying batch normalization (Ioffe and Szegedy, 2015), a further convolution and a residual (skip) connection (He et al.,

2016) to the beginning of the sub-block, and finally downsampling via average pooling. Our setup involved four such blocks, with four sub-blocks per block. The activation function used for all intermediate convolutional layers is the leaky rectified linear unit (LeakyReLU)

$$\text{LeakyReLU}(x) = \max(\alpha x, x).$$

After the encoder, the decoder

$$\mathcal{D} : \mathbb{R}^E \rightarrow \mathbb{R}^I$$

de-compresses \mathbf{l} into an approximation $\hat{\mathbf{x}}$ of \mathbf{x} . The decoder is similar to the encoder in structure, except the downsampling operations are replaced with upsampling operations via transpose convolutions.

In order to make the autoencoder useful for the flow reconstruction task, a further submodel

$$\mathcal{L} : \mathbb{R}^S \rightarrow \mathbb{R}^E$$

is necessary, which we call the latent space embedder. The latent space estimates latent space embeddings $\hat{\mathbf{l}}$ from the sensor inputs \mathbf{s} , which may be then used to compute $\hat{\mathbf{x}}$. It consists of a dense layer and several convolutional layers. The overall FR3D architecture is summarized in Fig. 5, and the parameter counts are displayed in Table 1.

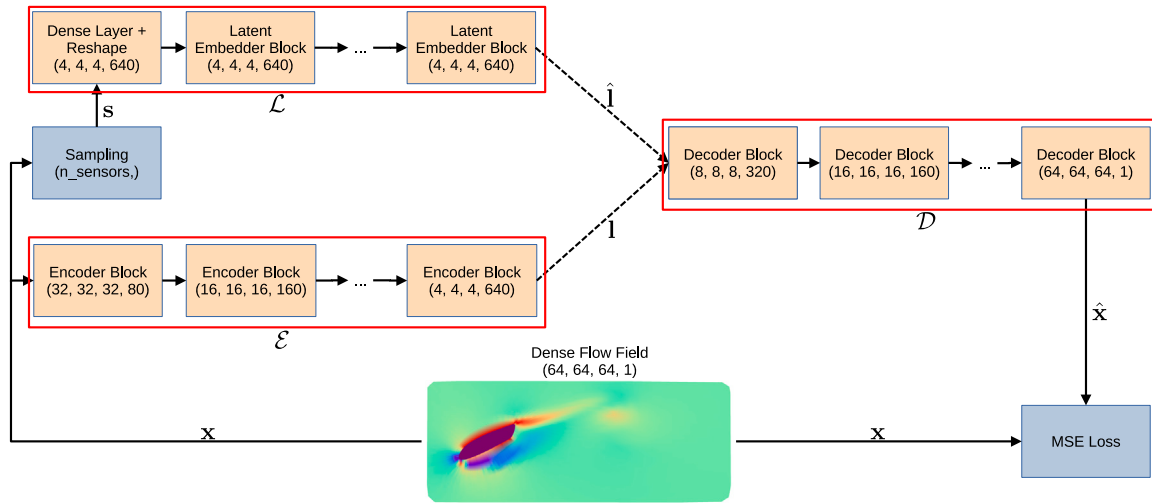


Fig. 5. Our model architecture for 3D flow reconstruction. The integer tuples in each block indicate the shape of the tensor outputted by that block. The first three dimensions correspond to the spatial dimensions, and the final dimension is the channel dimension.

Table 1

Parameter counts of the overall FR3D model and its submodels in the sparse sensor configuration.

	Submodel parameters	Total parameters
Encoder	124,910,960	402,163,801
Decoder	128,917,481	
Latent embedder	148,335,360	

The choice of architecture was motivated by two factors. First, a convolutional approach was adopted due to such architectures' strong relative performance in previous multi-geometry FR scenarios (Özbay and Laizet, 2022a). The autoencoder structure, meanwhile, was chosen to control the number of parameters in the fully connected layer ingesting the sensor inputs; many NN architectures used for FR incorporate a fully connected layer which accepts the low-dimensional sensor inputs and outputs a vector with dimensionality equal to the high-fidelity field. As the number of parameters of a fully connected layer scales linearly with the dimensionality of the output, and considering that the high-fidelity fields in 3D FR have a higher dimension compared to more commonly studied 2D FR scenarios, the autoencoder approach can substantially cut the number of parameters in the fully connected layers by constraining their outputs to the autoencoder's latent space.

The training of the autoencoder was conducted using the ADAM optimization algorithm with an initial learning rate of 10^{-4} , using the mean squared error (MSE) as the loss function L . An alternative approach using a generative adversarial network (GAN) (Goodfellow et al., 2020) instead of the more traditional MSE loss was also explored, however it was discarded in favor of the MSE loss function due to exhibiting lower accuracy, particularly for the v -velocity. Appendix provides more information about the GAN approach.

Optimization of the FR3D model using the MSE loss is done as follows: first, for each batch in the dataset, an optimization step is taken for the weights of \mathcal{E} and \mathcal{D} . Then, the weights of the decoder are 'frozen' and a further optimization step is taken for \mathcal{L} . This procedure, involving the dense field \mathbf{x} , the sensor inputs \mathbf{s} and the weights $\mathbf{w}_E, \mathbf{w}_D, \mathbf{w}_L$ of $\mathcal{E}, \mathcal{D}, \mathcal{L}$ respectively, is summarized in Algorithm 1.

The models and the training procedure were implemented using Tensorflow (Abadi et al., 2016) version 2.9 running on a server with two Nvidia A100 40 GB GPUs and an AMD EPYC 7443 24-core CPU. The number of training epochs was determined using an early stopping mechanism based on the validation loss, which automatically stops training when the validation loss level does not decline after a set number of epochs. To expedite the training process, once the weights of

Algorithm 1 Optimization of the FR3D model using a single batch of data

```

1: function OPTIMIZE( $\mathbf{w}_E, \mathbf{w}_D, \mathbf{w}_L, \mathbf{s}, \mathbf{x}$ )
2:    $\mathbf{l} := \mathcal{E}(\mathbf{x}, \mathbf{w}_E)$ 
3:    $\hat{\mathbf{x}} := \mathcal{D}(\mathbf{l}, \mathbf{w}_D)$ 
4:    $\mathbf{w}_E \leftarrow \text{ADAM}(\nabla_{\mathbf{w}_E} L(\mathbf{x}, \hat{\mathbf{x}}), \mathbf{w}_E)$ 
5:    $\mathbf{w}_D \leftarrow \text{ADAM}(\nabla_{\mathbf{w}_D} L(\mathbf{x}, \hat{\mathbf{x}}), \mathbf{w}_D)$ 
6:    $\hat{\mathbf{l}} := \mathcal{L}(\mathbf{s}, \mathbf{w}_L)$ 
7:    $\hat{\mathbf{x}} \leftarrow \mathcal{D}(\hat{\mathbf{l}}, \mathbf{w}_D)$ 
8:    $\mathbf{w}_L \leftarrow \text{ADAM}(\nabla_{\mathbf{w}_L} L(\mathbf{x}, \hat{\mathbf{x}}), \mathbf{w}_L)$ 
9: end function

```

\mathcal{E} and \mathcal{D} were obtained during training for the sparse sensor setup, they were reused for the plane sensor setup. The procedure was carried out separately for all flow variables (pressure p and velocity components u, v, w), each taking approximately 48 h to converge.

4. Results

Below, we present the results obtained by applying the FR3D model trained using the procedure in Section 3 to the validation dataset, consisting of 20 geometries not encountered during training. First, we display the results obtained for reconstructing the pressure and velocity from sparse sensors in Section 4.1, with Q-criterion contours also accurately reconstructed from the predicted velocity fields. Next, we demonstrate that the FR3D model can also be extended to estimate pressure fields from velocity data sampled on perpendicular planes in Section 4.2. Finally, we demonstrate that the pressure and velocity predictions from the FR3D model can be used to accurately estimate the time evolution of the drag and lift coefficients in Section 4.2.1.

4.1. Reconstruction from sparse sensors

We begin our analysis of the results using the sparse sensor case. Table 2 provides an overview of the model's overall performance via the mean absolute percentage error (MAPE) and mean squared error (MSE) metrics averaged over the entire validation dataset. To focus on the regions of highest interest in the domain, the error metrics are computed for sampling points inside a box with extents $[-L_m, 8L_m/3] \times [-L_m, L_m]$ in the x - and y -directions relative to the centroid of each object and covering the entire spanwise direction.

Table 2

Mean absolute percentage (MAPE) and mean squared (MSE) error levels achieved by the FR3D model on the validation dataset for reconstruction from sparse sensors. The “Input to D ” column refers to whether the model was run with the sensors (using the latent embedder submodel \mathcal{L} , inference configuration) or the ground truth field (using the encoder submodel \mathcal{E}) as the inputs.

Var.	Input to D	MAPE ^a	Min–max MAPE ^b	MSE	Min–max MSE
p	\mathcal{L}	9.50%	6.33%	6.21×10^{-3}	8.76×10^{-4}
	\mathcal{E}	3.68%	2.34%	1.86×10^{-3}	3.07×10^{-4}
u	\mathcal{L}	4.41%	2.68%	9.06×10^{-3}	2.20×10^{-4}
	\mathcal{E}	2.59%	1.36%	2.69×10^{-3}	6.37×10^{-5}
v	\mathcal{L}	16.33%	3.07%	6.31×10^{-3}	1.99×10^{-4}
	\mathcal{E}	11.96%	1.88%	2.08×10^{-3}	6.52×10^{-5}
w	\mathcal{L}	35.06%	3.52%	3.82×10^{-3}	5.15×10^{-4}
	\mathcal{E}	33.75%	2.69%	2.21×10^{-3}	2.97×10^{-4}

^aPercentage error figures are filtered to remove gridpoints with ground truth values less than 2% of the maximum absolute ground truth value in a snapshot.

^b“Min–max” refers to error figures with min–max normalization applied based on the ground truth field.

The error figures show that, overall, our model generalizes well to the validation set. The raw MAPE values for p and u are both below 10% for previously unseen geometries, which is in line with our previous work with 2D geometries (Özbay and Laizet, 2022b,a), despite the substantially greater challenge of 3D flows which contain more complicated structures orientated in various directions. In terms of absolute errors and normalized percentage errors, the predictions for v and w are also at a similar level of accuracy. However, we draw attention to the large discrepancy between MAPE and Min–max MAPE figures for v and w . This is caused by the fact that, due to our choice of boundary conditions, the mean ground truth values for p and u are distributed around a mean of 1, while those of v and w are distributed around 0. Due to this, though the absolute error levels (i.e. the numerator of the percentage error expression) are broadly similar for all four variables, the percentage error metrics for v and w are much higher since the denominator of the percentage error expression is much smaller for those variables.

Further evidence displaying the FR3D model’s generalization performance can be seen in Fig. 6, which visualizes the (encoder-computed) latent space vectors associated with validation snapshots via the t-SNE (Hinton and Roweis, 2002) method. The latent space vectors show that FR3D’s latent space is able to cluster the snapshots associated with different geometries together even for geometries unseen during training, which suggests that the latent space is well-conditioned for multi-geometry flow reconstruction.

Figs. 7, 8 and 9 provide further qualitative insights into the strong performance of the model via iso-contours of the Q-criterion, iso-contours of the pressure, and slices of both pressure and velocity components (respectively), for three randomly chosen snapshots exhibiting varying degrees of spanwise effects. 3D visualizations of these quantities from experimental data are often difficult as previously discussed – e.g. the amplification of sensor noise when computing velocity gradients presents a serious challenge for plotting the Q-criterion from experimental measurements. Thus, our results lay the groundwork for a step-function improvement in visualization of results from fluid dynamics experiments via flow reconstruction, all without the need for complicated post-processing techniques.

The contour plots indicate that the model is able to reconstruct the key details of the flow field. First, we draw attention to the shed vortices downstream of the object. In the top snapshot, the spanwise structure is placed at approximately $(2.1L_m, -0.2L_m)$ on the xy -plane; its shape is elliptic with a width of $0.2L_m$ and a height of $0.4L_m$. The middle snapshot shows two spanwise structures, one far downstream at $(2.2L_m, -0.1L_m)$ with a diameter of $0.5L_m$, and a second newly forming one at $(L_m, -0.1L_m)$. Finally, the bottom snapshot shows three structures, placed at $(0.8L_m, -0.1L_m)$, $(1.3L_m, 0.2L_m)$ and $(2.6L_m, -0.2L_m)$.




Fig. 6. t-SNE visualization of the FR3D model’s latent space colored by geometry, generated using the validation data. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The first and the third are roughly circular with $0.25L_m$ diameters, while the middle one is elliptical with a horizontal axis of $0.2L_m$ and a vertical axis of $0.4L_m$. Furthermore, these structures are slightly “bent” about the middle of the domain ($z = 1.5L_m$) in the x - and z - directions respectively by about $0.4L_m$. The locations and sizes of these structures are correctly reconstructed in all three cases and major features such as the “bends” in the bottom snapshot are present, although minor inaccuracies are present such as the thin section of the bend in the furthest downstream structure in the bottom snapshot.

Larger scale streamwise structures are also accurately reconstructed, particularly the hairpin structure spanning the box $[L_m, 2.2L_m] \times [-0.3L_m, 0.5L_m] \times [2.1L_m, 2.8L_m]$ in the first snapshot and the finger-like structures with size $1.2L_m$ aligned across the x -axis placed at $(0.5L_m, 1.5L_m)$ and $(0.5L_m, 2.5L_m)$ on the yz -plane connecting the two shed vortices in the second snapshot. Hence, our model is capable of replicating flows with different intensities of spanwise effects, which manifest as streamwise aligned vortices, spanwise aligned vortices, or a combination of both.

Pressure contours in Fig. 8 corroborate the observations from Q-criterion contours; the first two pressure snapshots display that the model is capable of reconstructing three dimensional features (the chimney-like structure in the first snapshot and the hole in the second snapshot) of the pressure field, while the third snapshot shows that the location and intensity of the shed vortices are correctly replicated. However, two areas of improvement stand out in both contour plots: first, compared to the ground truth, the predicted surfaces are noisier and less smooth. Second, some flow features at smaller length scales are missing; this is pronounced especially for the first snapshot in the figures.

Complementing the contour plots, Fig. 9 provides greater detail regarding the middle snapshot in Figs. 7 and 8, which is depicting a spanwise vortex that has just been shed from the trailing edge of the object, connected via streamwise structures to a second spanwise vortex further downstream. The core of the newly shed vortex in the snapshot is reconstructed clearly as shown by the xy -planes of the p and u plots, identifiable by the low pressure region in the p plot and the recirculation region in the u plot. Meanwhile, the yz -planes of the v and w provide further evidence of the high quality of the reconstructed streamwise structures; the negative v regions and the rapidly




Fig. 7. Ground truth (left) and predicted (right) Q -criterion ($Q = 0.05$) iso-contours from three randomly chosen snapshots belonging to different geometries.

alternating w regions are accurately reconstructed and correspond to the finger-like structures observed in Fig. 7. Fig. 10 depicts 2D maps of the spanwise-averaged absolute error for the predicted fields. Errors are chiefly coincident with the coherent structures in the flow, and do not increase with distance from sensor locations, highlighting the capability of the FR3D architecture to predict the flow accurately in the entire problem domain.

4.2. Reconstruction from plane measurements

Next, we showcase the results using the plane sensor setup. Table 3 displays error metrics with this sensor setup, similar to Table 2 for the sparse sensor setup. The error levels are slightly higher overall with this sensor setup, with the largest relative rise in error encountered in

predictions of u . Overall, the similarity between the error levels of the plane setup and the error levels with the sparse setup demonstrates the flexibility of our model in regards to the sensor setup, which is highly important for its future potential applications since its main envisioned future use – physical experiments – can involve widely varying sensor setups.

Pressure contours for three further geometries, using predictions made for this sensor setup, can be found in Fig. 11. Similar to the pressure contour results for the sparse sensor setup in Fig. 8, major features of the pressure field are accurately reconstructed with this sensor setup as well. Recovering pressure from plane measurements of velocity is often a challenge in experimental settings involving PIV (De Kat and Van Oudheusden, 2012). Thus, the present results suggest the FR3D model has the potential to help overcome challenges associated with recovering pressure fields from PIV experiments.

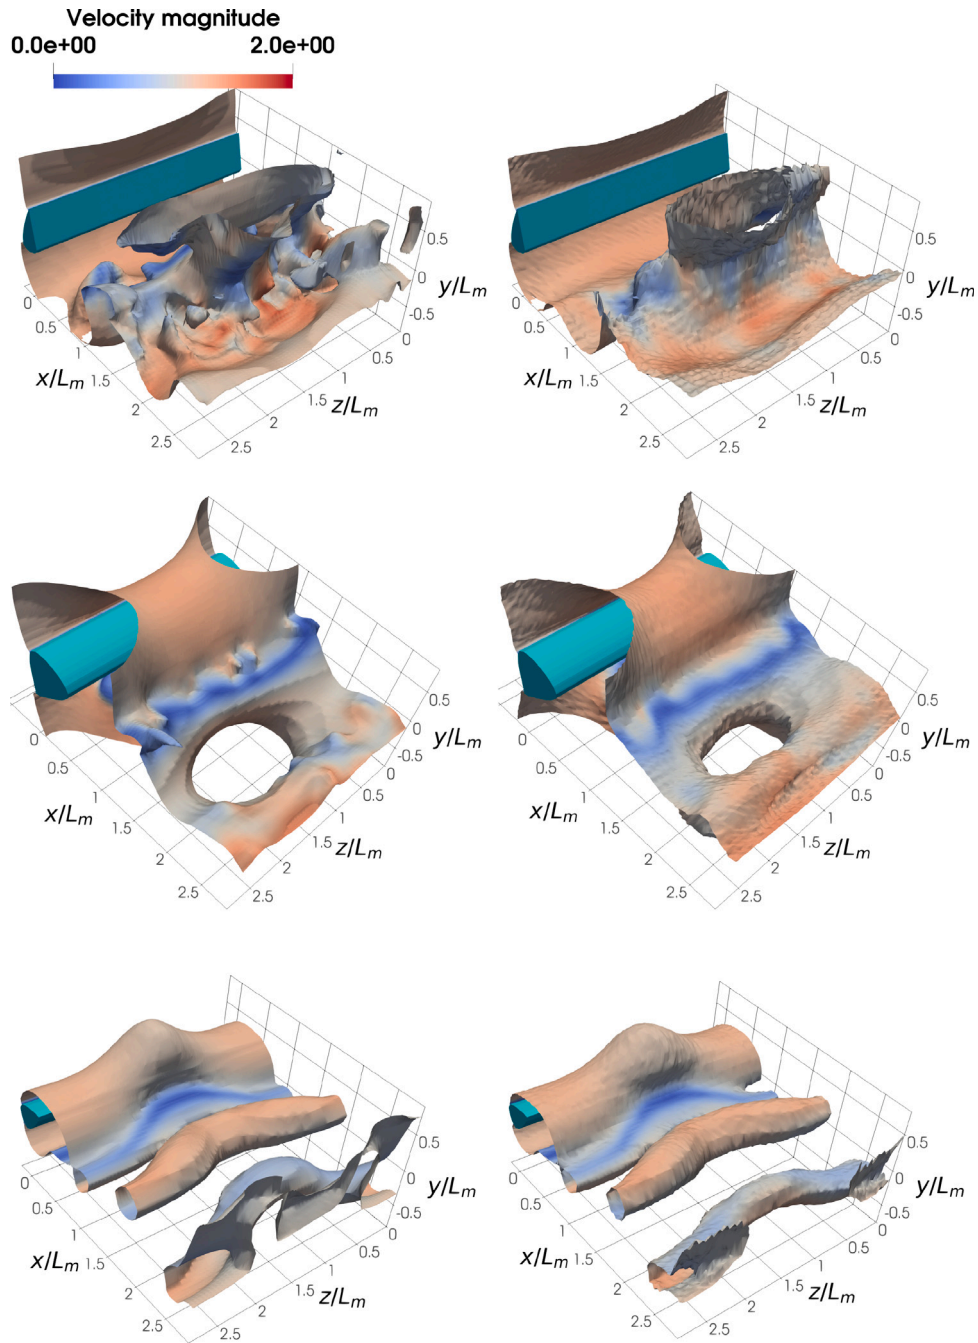


Fig. 8. Ground truth (left) and predicted (right) pressure (top, bottom: $p = 0.85$, middle: $p = 0.75$) iso-contours colored by velocity magnitude from three randomly chosen snapshots belonging to different geometries, using the model trained on the sparse sensor setup. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 3

Mean absolute percentage (MAPE) and mean squared (MSE) error levels achieved by the FR3D model on the validation dataset for reconstruction from plane velocity measurements. Error metrics using the encoder input are not provided, as they are identical to the values in Table 2.

Var.	Input to D	MAPE	Min–max MAPE	MSE	Min–max MSE
p	\mathcal{L}	10.19%	6.55%	6.55×10^{-3}	9.03×10^{-4}
u	\mathcal{L}	7.76%	5.17%	1.34×10^{-2}	9.66×10^{-4}
v	\mathcal{L}	18.15%	3.08%	6.96×10^{-3}	2.00×10^{-4}
w	\mathcal{L}	38.24%	5.77%	4.16×10^{-3}	7.43×10^{-4}

4.2.1. Estimation of lift and drag

Since the FR3D model is capable of accurately reconstructing the pressure and velocity fields, these results can be also applied to estimate the instantaneous lift and drag coefficients C_L and C_D experienced by the investigated geometries. Our choice of sampling points, utilizing the aforementioned conformal mapping approach instead of the more traditional Cartesian approach, makes this task substantially easier, as the need for interpolating the pressure and velocity fields onto the object surfaces is removed. This is due to the fact that the sampling points which lie on the inner ring of the annulus (cf. Fig. 3) always lie on the original geometry’s surface when the conformal mapping is applied.




Fig. 9. Ground truth (left) and predicted (right) slices of pressure and velocity fields for the middle snapshot in Figs. 7 and 8.




Fig. 10. Spanwise-averaged absolute error fields for the pressure, u -velocity, w -velocity and v -velocity (clockwise from top left), normalized by the standard deviations of the respective field variables, for the middle snapshot in Figs. 7 and 8.

To compute lift and drag forces, we adopt the approach of computing the body forces (pressure force and skin friction) through the integration of the pressure and the shear stresses across the object surface. The integration is carried out through the finite element method, approximating the surface of the geometry as a collection of quadrilateral surfaces. The dense field sampling points which lie on the object surface serve as the vertices of each quadrilateral. Since we do a simple extrusion in the z -direction, each quadrilateral is a rectangle. Defining a Cartesian coordinate system (ξ, η) on the quadrilateral, with the origin on one of the vertices of the quadrilateral, the standard bilinear basis functions

$$N_1 = (1 - \xi)(1 - \eta), \quad N_2 = \xi(1 - \eta), \\ N_3 = (1 - \xi)\eta, \quad N_4 = \xi\eta,$$

can be used to approximate the pressure distribution on quadrilateral j as

$$\tilde{p}_j = \sum_i \hat{p}_{ij} N_i,$$

where \hat{p}_{ij} is the pressure prediction on vertex i of quadrilateral j . Consequently, the pressure force on the quadrilateral (with surface normal $\hat{\mathbf{n}}_j$) can be approximated by integrating \tilde{p}_j along the surface of the quadrilateral:

$$\vec{f}_{p,j} = \int \tilde{p}_j \hat{\mathbf{n}}_j \, dA. \quad (1)$$

For a rectangle, the result of Eq. (1) is simply the average of the pressure values at its vertices, times its area

$$\vec{f}_{p,j} = \hat{\mathbf{n}}_j A_j \sum_{i=1}^4 \frac{1}{4} \hat{p}_{ij} = \vec{A}_j \sum_{i=1}^4 \frac{1}{4} \hat{p}_{ij}. \quad (2)$$

In practice, \vec{A}_j can be directly computed by taking the cross product of the vectors which run along the edges of the rectangle, and these vectors can be easily computed given the rectangle's vertex coordinates. The overall pressure force can be computed by adding the pressure force on each quadrilateral $\vec{f}_p = \sum_j \vec{f}_{p,j}$.

Table 4

Mean absolute percentage error (MAPE) and mean squared error (MSE) metrics averaged over the test dataset.

Sensor setup	Coefficient	MAPE	MSE
Sparse	C_L	9.16%	9.91×10^{-4}
	C_D	4.31%	4.86×10^{-4}
Planes	C_L	7.18%	6.91×10^{-4}
	C_D	3.43%	2.77×10^{-4}

The skin friction force \vec{f}_f can be computed by repeating the above procedure, substituting the wall shear stresses for the pressures. The shear forces can be approximated using the standard formula

$$\mu \frac{\partial u_{\parallel}}{\partial n},$$

where μ is the dynamic viscosity, u_{\parallel} is the velocity parallel to the surface and n is the coordinate normal to the surface. The gradient term can be approximated as

$$\frac{\partial u_{\parallel}}{\partial n} \approx \frac{u_{\parallel}(n = \delta n) - u_{\parallel}(n = 0)}{\delta n} = \frac{u_{\parallel}(n = \delta n)}{\delta n} \quad (3)$$

The annular sampling method also greatly simplifies the computation of the approximation in Eq. (3), as it ensures that each sampling point on the object surface (slice index 0 across the first dimension of the $64 \times 64 \times 64$ prediction array) has a corresponding sampling point above it in the wall-normal direction (slice index 1 across the first dimension of the same array).

Once both the skin friction forces \vec{f}_f and pressure forces \vec{f}_p are known, the overall body forces can be simply computed as $\vec{f} = \vec{f}_f + \vec{f}_p$. This method, combined with the FR3D model, leads to very accurate lift and drag predictions as seen in the error metrics averaged over the entire test dataset in Table 4 and the time evolutions of C_L and C_D plotted for two of the random test geometries in Fig. 12.

The lift and drag estimation capability presented by the FR3D model presents a substantial leap over similar approaches in previous




Fig. 11. Ground truth (left) and predicted (right) pressure (top: $p = 0.75$; middle, bottom: $p = 0.85$) iso-contours colored by velocity magnitude from three randomly chosen snapshots belonging to different objects, using the model trained on the plane sensor setup. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

literature; for example, [Chen et al. \(2021\)](#) demonstrated the capability to predict C_D with percentage errors of 3.43%, but only for steady 2D flows at Reynolds number equal to 10. In comparison, our model is capable of achieving similar levels of error, but for unsteady 3D flows for Reynolds numbers 50 times larger. The results are especially impressive when using the plane sensor setup, where we match the drag error levels reported by [Chen et al. \(2021\)](#), despite the substantially more difficult task at hand.

We observed that our lift and drag predictions are especially accurate on high aspect ratio shapes resembling airfoils, such as the geometry on the right column in [Fig. 12](#). The level of accuracy on bluff bodies is still reasonably high, as seen in the bluff body on the left column in the same figure. However, we did see a slight degradation

in accuracy on shapes with concave sections, such as the middle right geometry in [Fig. 1](#) for which the time-averaged C_L and C_D prediction errors were 14% and 8% respectively. This is due to the substantially more complicated flow patterns occurring in such geometries, and the relative rarity of such geometries in the dataset.

5. Conclusion and future work

In this study, we described the performance of FR3D, a convolutional autoencoder neural network model, on the reconstruction of three-dimensional flows past objects with varying cross-sections, given measurements of the flow field from sparse sensors and plane measurements. Using a conformal mapping technique to achieve geometry

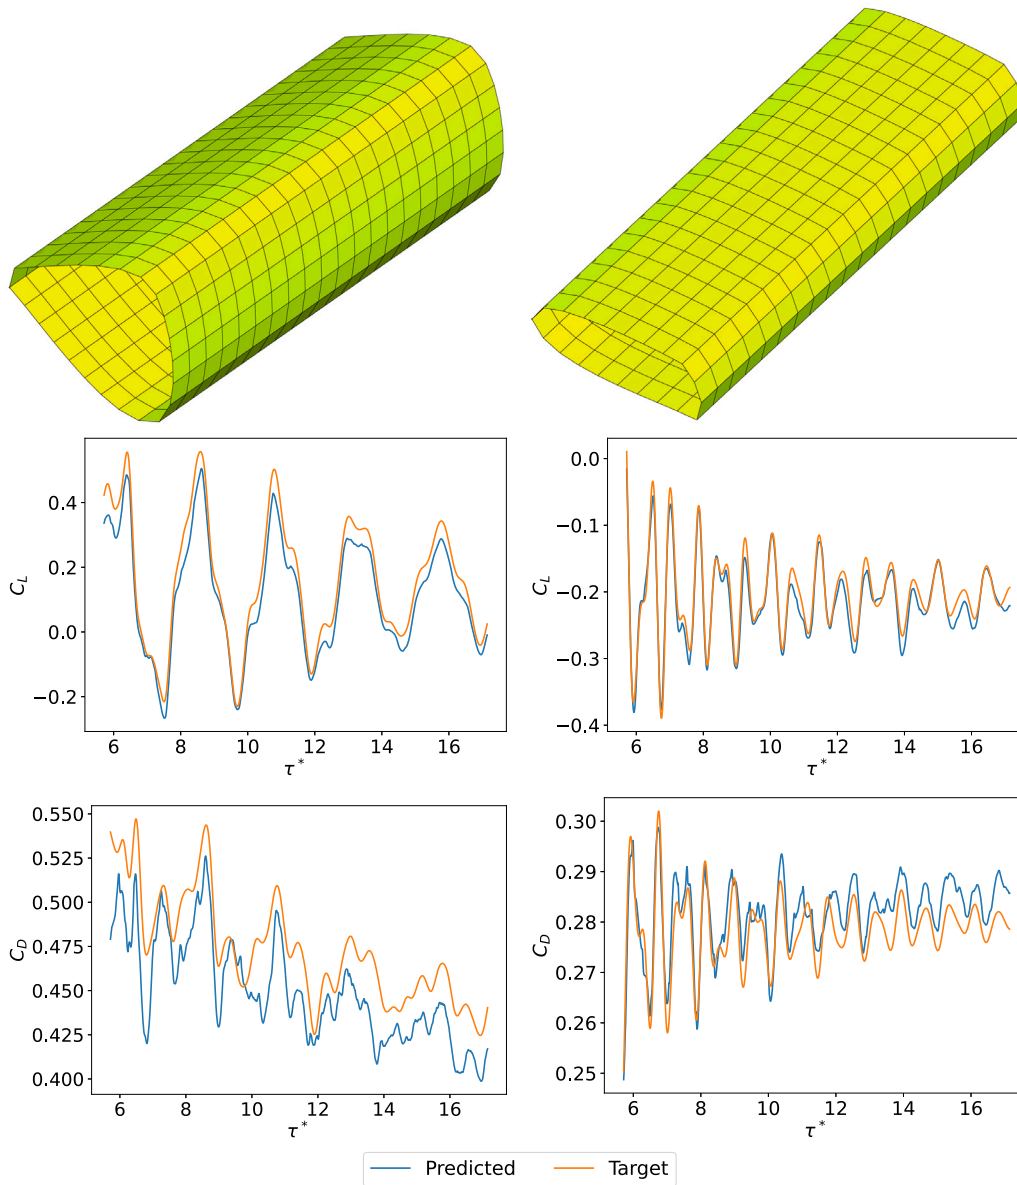


Fig. 12. Time evolutions of the lift (middle) and drag (bottom) coefficients for two of the random geometries (left and right) in the test dataset. Predictions were made using the model trained using the sparse sensor setup, and the target values were computed using the body force results outputted by PyFR.

invariance, we demonstrated that the FR3D architecture is capable of reconstructing instantaneous pressure and velocity fields of flows past such geometries with min–max normalized percentage error levels under 10% for geometries not encountered during training at a fixed Reynolds number. The reconstructions are of a quality sufficient to accurately replicate the major features of Q-criterion and pressure iso-contours.

Subsequently, we applied the FR3D model to a scenario in which velocity measurements are available in two orthogonal planes, whereby it is attempted to recover the flow variables – including pressure – from the measurements of the velocity components downstream of the object. The FR3D model also performed well in this additional scenario, reconstructing the dense pressure fields with percentage errors just above 10%.

Finally, using the reconstructed fields, we demonstrated that the lift and drag coefficients can be estimated within 10% of the ground truth values using both sensor setups, going as low as 3.43% when estimating drag coefficients using the plane measurement setup.

In the future, we aim to extend this work by investigating:

- **Noisy measurements:** The FR3D model was trained and evaluated using high-fidelity values obtained via computation. In contrast, a large degree of uncertainty exists in real-world measurements and ensuring that FR models are robust to noise will be necessary before they can be utilized in real laboratory environments, as opposed to computational studies.
- **More complex geometries:** The present work focused on investigating only objects extruded in the spanwise direction, in order to leverage our previous work on applying Schwarz–Christoffel mappings for training a model that can handle different geometries well. Different mapping techniques, such as boundary conforming curvilinear coordinate systems (Thompson, 1982), or different neural network model architectures such as graph neural networks that do not depend on regular grids, will be necessary to achieve geometry invariance for a broader class of objects.
- **Varying Reynolds numbers:** Since the dataset in this work consisted solely of flows at a fixed Reynolds number, the model

is not expected to perform well at other Reynolds numbers. Extending the model to perform well for a wide range of Reynolds numbers, possibly through adding a ‘physics-informed’ loss function component, will greatly boost its usefulness. Reconstructing three-dimensional flows at high Reynolds numbers, relevant to real-life applications might be possible with hybrid approaches, combining different neural network architectures. For example, combining CNNs with recurrent neural networks (RNNs) could better capture both spatial and temporal dependencies in the flow field.

- **State-of-the-art generative models:** Recent advances in generative models for images, such as Diffusion architectures (Rombach et al., 2022), constitute promising directions for substantial advances in flow reconstruction.

CRedit authorship contribution statement

Ali Girayhan Özbay: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing – original draft, Visualization. **Sylvain Laizet:** Conceptualization, Resources, Writing – review & editing, Supervision, Funding acquisition.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Ali Girayhan Özbay reports equipment, drugs, or supplies was provided by NVIDIA Corp.

Data availability

Github link in manuscript.

Acknowledgments

This work was supported by a PhD scholarship by the Department of Aeronautics, Imperial College London, United Kingdom, and an Academic Hardware Grant by NVIDIA, United States. The code used to generate the data and train the models used in this study is available at <https://github.com/aligirayhanozbay/fr3D>

Appendix. Results using the GAN approach

Flow reconstruction models, such as the FR3D model introduced in the present work, can be broadly described as ‘generative’ deep learning models. Generative adversarial networks (GANs) (Goodfellow et al., 2020) constitute a method of training generative deep learning models by introducing a second neural network called a ‘discriminator’. The discriminator is tasked with classifying whether a particular output belongs to the ground truth dataset or was generated by the generator. Throughout the training process, the generator tries to ‘fool’ the discriminator; the discriminator outputs probabilities regarding whether the outputs of the generator were generated or not, and the generator tries to minimize that probability by using the discriminator as a component of its loss function.

Training the FR3D model with a GAN approach was attempted using a procedure based on Algorithm 1. The new approach first trains the FR3D encoder \mathcal{E} , decoder D and latent space embedder \mathcal{L} in a similar fashion, but adds a binary cross-entropy loss component

$$L_c(\mathbf{t}, \mathbf{y}) = -(\mathbf{y} \ln(\mathbf{t}) + (1 - \mathbf{y}) \ln(1 - \mathbf{t})), \quad (\text{A.1})$$

for the outputs of the discriminator C . The outputs from the decoder are then used to train the discriminator with a batch of data consisting of both the decoder outputs and the ground truth snapshots. This procedure is outlined in Algorithm 2.

Algorithm 2 Optimization of the GAN variant of the FR3D model using a single batch of data. $\text{ZERONES}(\cdot)$ and $\text{ONES}(\cdot)$ are functions which output a vector of 0 or 1 values, respectively, with the number of entries equal to the number of samples in the input. $\text{CONCAT}(\cdot, \cdot)$ is a function which concatenates two tensors along their first dimension.

```

1: function OPTIMIZEGAN( $\mathbf{w}_E, \mathbf{w}_D, \mathbf{w}_L, \mathbf{w}_C, \mathbf{s}, \mathbf{x}$ )
2:    $\mathbf{t}_D := \text{ZERONES}(\mathbf{x})$ 
3:    $\mathbf{l} := \mathcal{E}(\mathbf{x}, \mathbf{w}_E)$ 
4:    $\hat{\mathbf{x}} := D(\mathbf{l}, \mathbf{w}_D)$ 
5:    $\mathbf{w}_E \leftarrow \text{ADAM}(\nabla_{\mathbf{w}_E} L(\mathbf{x}, \hat{\mathbf{x}}) + L_c(\mathbf{t}_D, C(\hat{\mathbf{x}}, \mathbf{w}_C)), \mathbf{w}_E)$ 
6:    $\mathbf{w}_D \leftarrow \text{ADAM}(\nabla_{\mathbf{w}_D} L(\mathbf{x}, \hat{\mathbf{x}}) + L_c(\mathbf{t}_D, C(\hat{\mathbf{x}}, \mathbf{w}_C)), \mathbf{w}_D)$ 
7:    $\hat{\mathbf{l}} := \mathcal{L}(\mathbf{s}, \mathbf{w}_L)$ 
8:    $\hat{\mathbf{x}} \leftarrow D(\hat{\mathbf{l}}, \mathbf{w}_D)$ 
9:    $\mathbf{w}_L \leftarrow \text{ADAM}(\nabla_{\mathbf{w}_L} L(\mathbf{x}, \hat{\mathbf{x}}) + L_c(\mathbf{t}_D, C(\hat{\mathbf{x}}, \mathbf{w}_C)), \mathbf{w}_L)$ 
10:   $\mathbf{t}_C := \text{CONCAT}(\text{ONES}(\mathbf{x}), \text{ZERONES}(\mathbf{x}))$ 
11:   $\mathbf{w}_C \leftarrow \text{ADAM}(\nabla_{\mathbf{w}_C} L_c(\mathbf{t}_C, C(\text{CONCAT}(\hat{\mathbf{x}}, \mathbf{x}), \mathbf{w}_C)), \mathbf{w}_C)$ 
12: end function

```

Table A.5

Mean absolute percentage (MAPE) and mean squared (MSE) error levels achieved by the FR3D model trained using the GAN approach on the validation dataset for reconstruction from point measurements.

Var.	Input to D	MAPE	Min–max MAPE	MSE	Min–max MSE
p	\mathcal{L}	11.15%	7.74%	6.99×10^{-3}	1.05×10^{-3}
u	\mathcal{L}	10.98%	8.12%	8.21×10^{-3}	1.57×10^{-3}
v	\mathcal{L}	44.06%	15.08%	1.66×10^{-2}	3.61×10^{-3}
w	\mathcal{L}	35.16%	7.29%	3.29×10^{-3}	1.32×10^{-3}

The NN architecture chosen for the discriminator C was identical to the encoder, however the pool size was increased to 4 and a flatten operation followed by a dense layer using the sigmoid activation function was appended to classify the inputs. The results from training the FR3D model with this GAN setup are summarized in Table A.5.

The GAN approach achieves largely similar error levels to MSE for the pressure field and w , however also exhibits severe degradation in performance for u and v . As a result, it was observed that the FR3D model trained with the GAN approach had substantially worse performance than the baseline MSE version in terms of aerodynamic force coefficient predictions, with the mean percentage errors for C_L and C_D increasing to 23.39% and 14.29% respectively. Predictions of the location and intensity of vortices/coherent structures in the wakes are also similarly worse, with the Q-criterion contours not displaying many of the features otherwise reconstructed in the MSE version.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al., 2016. Tensorflow: A system for large-scale machine learning. In: 12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16). pp. 265–283.
- Akeweje, E., Avilkin, V., Olhin, A., Panov, M., Vishnyakov, A., 2023. Real-Time Reconstruction of Complex Flow in Nanoporous Media: Linear vs Non-linear Decoding.
- Bonnet, J.P., Cole, D.R., Delville, J., Glauser, M.N., Ukeiley, L.S., 1994. Stochastic estimation and proper orthogonal decomposition: complementary techniques for identifying structure. Exp. Fluids 17 (5), 307–314.
- Borée, J., 2003. Extended proper orthogonal decomposition: a tool to analyse correlated events in turbulent flows. Exp. Fluids 35 (2), 188–192.
- Callahan, J.L., Maeda, K., Brunton, S.L., 2019. Robust flow reconstruction from limited measurements via sparse representation. Phys. Rev. Fluids 4, 103907.
- Carter, D.W., De Voogt, F., Soares, R., Ganapathisubramani, B., 2021. Data-driven sparse reconstruction of flow over a stalled aerofoil using experimental data. Data-Cent. Eng. 2, e5. <http://dx.doi.org/10.1017/dce.2021.5>.
- Chandramouli, P., Memin, E., Heitz, D., Fiabane, L., 2019. Fast 3D flow reconstructions from 2D cross-plane observations. Exp. Fluids 60, 1–21.
- Chen, J., Hachem, E., Viquerat, J., 2021. Graph neural networks for laminar flow prediction around random two-dimensional shapes. Phys. Fluids 33 (12), 123607, [arXiv:https://doi.org/10.1063/5.0064108](https://doi.org/10.1063/5.0064108).

- De Kat, R., Van Oudheusden, B., 2012. Instantaneous planar pressure determination from piv in turbulent flow. *Exp. Fluids* 52 (5), 1089–1106.
- Dubois, P., Gomez, T., Planckaert, L., Perret, L., 2022. Machine learning for fluid flow reconstruction from limited measurements. *J. Comput. Phys.* 448, 110733.
- Duthé, G., Abdallah, I., Barber, S., Chatzi, E., 2023. Graph neural networks for aerodynamic flow reconstruction from sparse sensing.
- Erichson, N.B., Mathelin, L., Yao, Z., Brunton, S.L., Mahoney, M.W., Kutz, J.N., 2020. Shallow neural networks for fluid flow reconstruction with limited sensors. *Proc. R. Soc. A* 476 (2238), 20200097, arXiv:https://royalsocietypublishing.org/doi/pdf/10.1098/rspa.2020.0097.
- Fukami, K., Fukagata, K., Taira, K., 2020. Assessment of supervised machine learning methods for fluid flows. *Theor. Comput. Fluid Dyn.* 34, 497–519.
- Fukami, K., Maulik, R., Ramachandra, N., Fukagata, K., Taira, K., 2021. Global field reconstruction from sparse sensors with voronoi tessellation-assisted deep learning. *Nat. Mach. Intell.* 3 (11), 945–951.
- Funahashi, K.-I., 1989. On the approximate realization of continuous mappings by neural networks. *Neural Netw.* 2 (3), 183–192. [http://dx.doi.org/10.1016/0893-6080\(89\)90003-8](http://dx.doi.org/10.1016/0893-6080(89)90003-8), URL <https://www.sciencedirect.com/science/article/pii/0893608089900038>.
- Giannenas, A.E., Laizet, S., 2021. A simple and scalable immersed boundary method for high-fidelity simulations of fixed and moving objects on a cartesian mesh. *Appl. Math. Model.* 99, 606–627.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2020. Generative adversarial networks. *Commun. ACM* 63 (11), 139–144.
- Guastoni, L., Güemes, A., Ianiro, A., Discetti, S., Schlatter, P., Azizpour, H., Vinuesa, R., 2021. Convolutional-network models to predict wall-bounded turbulence from wall quantities. *J. Fluid Mech.* 928, A27.
- Güemes, A., Discetti, S., Ianiro, A., Sirmacek, B., Azizpour, H., Vinuesa, R., 2021. From coarse wall measurements to turbulent velocity fields through deep learning. *Phys. Fluids* 33 (7), 075121.
- He, X., Wang, Y., Li, J., 2022. Flow completion network: inferring the fluid dynamics from incomplete flow information using graph neural networks. *Phys. Fluids* 34 (8), 087114. <http://dx.doi.org/10.1063/5.0097688>, arXiv:https://doi.org/10.1063/5.0097688.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: 2016 IEEE conference on computer vision and pattern recognition (CVPR). pp. 770–778.
- Hinton, G.E., Roweis, S., 2002. Stochastic neighbor embedding. In: Becker, S., Thrun, S., Obermayer, K. (Eds.), *Advances in Neural Information Processing Systems*. vol. 15, MIT Press.
- Huynh, H.T., 2007. A flux reconstruction approach to high-order schemes including discontinuous galerkin methods. In: 18th AIAA computational fluid dynamics conference. p. 4079.
- Ioffe, S., Szegedy, C., 2015. Batch normalization: accelerating deep network training by reducing internal covariate shift. In: *International conference on machine learning*. PMLR, pp. 448–456.
- Jeong, J., Hussain, F., 1995. On the identification of a vortex. *J. Fluid Mech.* 285, 69–94.
- Kumar, Y., Bahl, P., Chakraborty, S., 2021. State estimation with limited sensors—A deep learning based approach. <https://www.sciencedirect.com/science/article/pii/S0021999122001437>.
- Laima, S., Zhou, X., Jin, X., Gao, D., Li, H., 2023. DeepTRNet: Time-resolved reconstruction of flow around a circular cylinder via spatiotemporal deep neural networks. *Phys. Fluids* 35 (1), 015118.
- Liu, Q., Zhu, W., Ma, F., Jia, X., Gao, Y., Wen, J., 2022. Graph attention network-based fluid simulation model. *AIP Adv.* 12 (9), 095114.
- Matsuo, M., Nakamura, T., Morimoto, M., Fukami, K., Fukagata, K., 2021. Supervised convolutional network for three-dimensional fluid data reconstruction from sectional flow fields with adaptive super-resolution assistance.
- Mittal, R., Balachandar, S., et al., 1997. On the inclusion of three-dimensional effects in simulations of two-dimensional bluff-body wake flows. In: *ASME fluids engineering division summer meeting*. Citeseer, pp. 1–6.
- Nakamura, T., Fukagata, K., 2022. Robust training approach of neural networks for fluid flow state estimations. *Int. J. Heat Fluid Flow* 96, 108997.
- Özbay, A.G., Laizet, S., 2022a. Deep learning fluid flow reconstruction around arbitrary two-dimensional objects from sparse sensors using conformal mappings. *AIP Adv.* 12 (4), 045126.
- Özbay, A.G., Laizet, S., 2022b. Unsteady two-dimensional flow reconstruction and force coefficient estimation around arbitrary shapes via conformal mapping aided deep neural networks. In: *Proceedings of TSFP-12*. no. 195.
- Pérez, J.M., Le Clairche, S., Vega, J.M., 2020. Reconstruction of three-dimensional flow fields from two-dimensional data. *J. Comput. Phys.* 407, 109239.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B., 2022. High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695.
- Schmid, P.J., 2010. Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* 656, 5–28.
- Schweri, L., Foucher, S., Tang, J., Azevedo, V.C., Günther, T., Solenthaler, B., 2021. A physics-aware neural network approach for flow data reconstruction from satellite observations. *Front. Clim.* 3.
- Sun, L., Wang, J.-X., 2020. Physics-constrained bayesian neural network for fluid flow reconstruction with sparse and noisy data. *Theor. Appl. Mech. Lett.* 10 (3), 161–169. <http://dx.doi.org/10.1016/j.taml.2020.01.031>, URL <https://www.sciencedirect.com/science/article/pii/S2095034920300295>.
- Taylor, J.A., Glauser, M.N., 2004. Towards practical flow sensing and control via POD and LSE based low-dimensional tools. *J. Fluids Eng.* 126 (3), 337–345.
- Thompson, J.F., 1982. General curvilinear coordinate systems. *Appl. Math. Comput.* 10, 1–30.
- Verma, S., Novati, G., Koumoutsakos, P., 2018. Efficient collective swimming by harnessing vortices through deep reinforcement learning. *Proc. Natl. Acad. Sci.* 115 (23), 5849–5854.
- Viquerat, J., Hachem, E., 2020. A supervised neural network for drag prediction of arbitrary 2d shapes in laminar flows at low reynolds number. *Comput. & Fluids* 210, 104645.
- Willcox, K., 2006. Unsteady flow sensing and estimation via the gappy proper orthogonal decomposition. *Comput. & Fluids* 35 (2), 208–226.
- Witherden, F., Farrington, A., Vincent, P., 2014. PyFR: An open source framework for solving advection–diffusion type problems on streaming architectures using the flux reconstruction approach. *Comput. Phys. Comm.* 185 (11), 3028–3040.
- Xu, S., Sun, Z., Huang, R., Guo, D., Yang, G., Ju, S., 2023. A practical approach to flow field reconstruction with sparse or incomplete data through physics informed neural network. *Acta Mechanica Sinica* 322302, <https://link.springer.com/article/10.1007/s10409-022-22302-x>.